

# Internet Mini-cores

*Local Communications in the Internet's "spur" regions*

*Steve Gibbard*

*[Packet Clearing House/Steve Gibbard Consulting](#)*

*Originally presented at [SANOG](#) in Dhaka, Bangladesh, February 13, 2005.*

Internet Mini-cores.....	1
Current Internet Structure .....	1
Implications.....	2
Examples.....	2
Proposed new model.....	3
How to get there.....	3
Progress.....	4
More needs to be done .....	5
Documentation required.....	5
Caveats .....	6
About the Author .....	6
Contact information .....	6

## ***Current Internet Structure***

The Internet currently consists of a well connected core and less well connected spurs. Within the core, connectivity is good. There is lots of fiber, lots of redundancy, and lots of cheap bandwidth. The ability to send large amounts of data quickly between urban areas of the "developed world" can now be taken for granted. For Internet users in the core region, life is good.

For the rest of the world, areas that relate to the core as long spurs, it's a different story. Many different ISPs in these regions have connectivity to the core, but few connect to each other. The connections to the core often extend hundreds or thousands of miles, sometimes via satellite links to other continents. Even "local" connectivity uses these connections, so data has to go half way around the world and back in order to go a few miles, or less. International connectivity in these regions is expensive, costing \$5,000 or more per megabit per second, while latency on a satellite link is around half a second. Since local connectivity crosses two satellite links, it ends up being twice as slow, and twice as expensive.

This also creates reliability problems. The longer a circuit is, the more there is that can go wrong with it, and this effect is amplified in situations where fiber is considered too difficult and expensive. While true local communications are vulnerable mainly to problems with fiber links of a few miles, long distance communications are vulnerable to failures of microwave and satellite links going long distances, often through the networks of multiple carriers. When local communications are carried on long distance

connections, local communications become vulnerable to the reliability problems of long distance circuits.

The purpose of this paper is to examine the problems this structure causes for the spur regions of the Internet, and to propose some solutions.

## **Implications**

With traditional phone networks, there is a big advantage to making local calls. Our definitions of local have shifted over time. With the deployment of plentiful fiber in many parts of the world, calls to places progressively farther and farther away have become “too cheap to meter.” But even from the best connected parts of the world, there are still some places which, due to distance, politics, or poor connectivity, it’s more expensive to call than places considered more local.

The Internet has changed the thinking on that, leading to some well publicized pronouncements that “distance is dead.” This has been widely touted as a feature, and in many cases it is, but the impact on areas outside the Internet core has not been to make all communications local, but to make all communications long distance. While rates for packet switched data are much lower than those for circuit switched connections, those communicating locally in those parts of the world end up sending their data to far away places, and paying for it. This drives up the cost of real long distance communications as well, by increasing demand for use of the long distance circuits.

Furthermore, local phone connectivity tends to work regardless of what’s happening in other parts of the world. Outages that affect local calls tend to be local. Since local connectivity tends to be cheap, it’s reasonably easy to provision enough redundancy to make such outages rare. Outages of long distance services may be more common, but few people notice when there’s a problem making international calls. On the Internet, in some parts of the world, all connectivity is international, with all of the vulnerabilities that entails.

As more and more phone service moves to the Internet, under the current model local phone calls would become international calls as well.

## **Examples**

In August, 2004, there was a fiber cut off the coast of Sri Lanka. A freighter had dropped anchor in the wrong place, and severed the fiber connectivity between Sri Lanka and the rest of the world. Media reports declared this as a loss of “Internet and long distance phone service, and continued on with the implication that Sri Lanka and the Internet were two separate things that connect to each other. In fact, what had happened was that various pieces of the Internet inside Sri Lanka had been disconnected from the Internet outside.

I suspect the story was not quite so simple. There certainly wasn't working IP communication between Sri Lanka and the rest of the world, but they do have an exchange point so there should have been local connectivity. However, lacking a root DNS server locally, their local Internet communications would still have been entirely dependent on access to the outside, which is not a recipe for reliability.

Cost considerations are also somewhat staggering. In Nepal, the cost of sending data locally through the Nepal Internet Exchange is around \$50 per Mb/s, while the cost of sending data internationally through the satellite connections is around \$5000 per Mb/s. Without the exchange point, all traffic gets billed as international. For more information about this, please see my "Economics of Peering" paper.

### ***Proposed new model***

There's nothing wrong with a global well-connected network like the current Internet core. It serves a lot of the world very well. Still, we shouldn't require everything to go through it, as this has some tremendous negative implications for those parts of the world that aren't well connected. A better model would be to go from the one Internet core that everybody now tries to get close to, to a model in which the Internet has many regional cores, all with good connectivity internally, and connectivity between them to be used as needed. As this would increase reliability and performance, while decreasing costs, it seems like a big win all around.

Local connectivity should be used for local communications – data sent to your neighbors shouldn't go through other continents. ISPs should still have access to connectivity out of their region, but that connectivity should be used for data that is actually leaving the region, not just being sent around the world to come right back. There are a number of currently underway efforts to make this happen, and a lot more that needs to be done.

### **How to get there**

The first step is to make sure local traffic stays local. For this, there should be a local exchange point, which all local ISPs should have connectivity to. By handing traffic to each other via the exchange point, they avoid needing to use expensive long distance links to connect to each other, and avoid being dependant on connectivity to the far away core. While all local networks need connectivity to the exchange point for this to be fully effective, not all networks need to connect to the exchange point directly. Buying transit from a network that is fully peered at the exchange point should be sufficient, as buying transit is effectively paying somebody else to do your peering for you.

The same result could be gotten through the use of a monopoly transit provider – all local data would go to the monopoly transit provider, who would hand it off to the other local ISPs without the traffic having to leave the region. However, this doesn't mean a

monopoly transit provider is as good as an exchange point. Monopolies have little incentive to provide good service, and monopolies tend to lead to high prices, defeating part of the purpose.

Building exchange points isn't enough. Full in-region connectivity doesn't help much when what you need to connect to is elsewhere. For example, connectivity at layer 3 doesn't do much when you're cut off from DNS, unless you know the IP addresses of the hosts you're trying to connect to. Few people do, for good reason. Even with DNS resolution, if what you're trying to connect to is elsewhere, say Hotmail or a SIP server in another region, it doesn't help that whoever you're trying to e-mail or call is local.

At a minimum, therefore, it is important for every internally-connected region to have an anycast copy of one of the root DNS servers, and locally hosted copies of any other domains in local use. This should presumably include at least the local country code TLD, and any second and subsequent level domains that are used to access hosts in the region. Of particular note gTLDs (.com, .net, .org, etc.) and ccTLDs for other countries should not be considered to be reliable in regions where those TLD operators don't have local servers.

Beyond the DNS, defining critical services gets a bit more subjective. E-mail seems obvious. E-mail servers operated by local ISPs already fill this need in many cases, but due to their popularity it would be nice if there were also local copies of Yahoo Mail and Hotmail, or local equivalents of them. As more phone service moves to the Internet, local SIP servers and VOIP to PSTN gateways also seem like an obvious requirement, along with the necessary branches of the Enum DNS tree to support locally used phone numbers.

Beyond that, the definition of critical services becomes more one of philosophy. For example, is Google a critical service? This is really a matter for content providers to consider. They need to determine where they want their services to work reliably, and place their servers accordingly.

## **Progress**

Lots of work is being done on making the Internet more stable in regions with poor outside connectivity, and many of the things mentioned in this paper are being done.

There have been a lot of efforts to build exchange points in developing countries, so at least some local traffic is now able to be kept local in Bangladesh, Kenya, Mozambique, Nepal, Sri Lanka, Tanzania, Uganda, and several other countries.

Local TLD operators are in many cases hosting at least one of their name servers in the region covered by the TLD. In regions where there is a local exchange point, the TLD will have local reachability if its servers are hosted at an ISP that connects to the exchange point, and in many cases this is happening. Hosting copies of their TLDs

outside of their regions is also desirable; the important thing is that it also be reachable locally.

Through the use of anycast, some of the root server operators have started getting root server copies distributed more widely, but in many cases this still appears to be mostly an effort to get the servers better distributed throughout the “developed world.” ISC does now have copies of the F root server in Jakarta and Johannesburg, but this is an issue on which a lot more needs to be done.

In many cases, local content providers are starting to host their content locally. Their users see the results of this in faster connections, and in some places being able to get to the websites without having to pay for international bandwidth.

## **More needs to be done**

As mentioned above, many regions still don’t have local exchange points, making them entirely dependent on their connections to the outside world for local communications. It would be a great benefit to the ISPs and Internet users in those regions to solve that problem. Without a local exchange, locally hosted services tend to be of little benefit. They end up topologically twice as far away from many users in their local region as they would be if they were hosted in the core.

What we’ve seen in many places is that once there’s a local exchange, other local services tend to follow. The exceptions to this so far have been root DNS and the gTLDs. In the case of the DNS root, this is likely a matter of the DNS root being controlled by a relatively small set of organizations, making it hard for root servers to be set up by a local grass roots effort. Some of the root server operators are now making local root server copies available to those able to come up with funding, but that funding is generally much easier to come by in areas that can also afford good connectivity to the rest of the world. The gTLDs, being operated under contract mostly by for-profit operators, may be a bigger problem – for-profits, especially for-profit monopolies, tend to go where expenses are low and revenue potential is high. It would be nice if ICANN would start including providing service in poorly connected areas as a requirement in renewing gTLD operations contracts, but until then Internet services in poorly connected regions would be best served by avoiding using the gTLDs, and using their local ccTLDs instead.

## **Documentation required**

After years of the “distance is dead” argument, Internet users haven’t been conditioned to think about the locations of the servers they are accessing. Those of us who have tried to explain backbone engineering jobs to non-technical people learn this fast – explaining that you work on getting data from servers in one place to users in another draws blank stares from people who don’t realize that the Web isn’t just something that magically

comes from their ISP. Convincing users to care where the services they rely on are located – that they're better off using something local than something that looks cool but is half way around the world – may take some considerable work. On the other hand, if there are significant performance differences between local and remote content, as there sometimes are, users may end up preferring the local content on their own.

For users who need help understanding this, it may help for the content providers and ISPs they are using to document where their services are most useful. The documentation from the content providers doesn't have to look like a disclaimer – “locally hosted in location X” could look like a feature to users in location X, leaving the competition with the burden of responding. The Internet service providers can help by guiding their users towards local content where possible, thus saving money on their long-distance transit and cutting down on customer complaints about speed and reliability.

## ***Caveats***

Despite what is said here about the value of local communication, this should not be seen as an attack on long-distance communication. Lowering the costs of long-distance communications is one of the great achievements of the Internet, and has done a lot to bring the world closer together. Still, long-distance communication does still suffer from some considerable cost constraints, and the improvements required for long-distance communication are well understood; just hard to afford. I'd like to free local communications from the constraints that affect long-distance communications.

## ***About the Author***

*[Steve Gibbard](#) is the Network Architect at [Packet Clearing House](#). He runs a global research network and an anycast DNS network that hosts the top level domains for several countries as well as a root DNS server, and studies the interconnection of Internet networks around the world. In addition, he does network architecture and peering work as a consultant for several ISPs in the Bay Area and elsewhere. Steve is a former Senior Network Engineer at Cable & Wireless, and has held network engineering positions at Digital Island and World Wide Net.*

## **Contact information**

Steve Gibbard

[scg@pch.net](mailto:scg@pch.net)

<http://www.pch.net>

<http://www.stevegibbard.com>

+1 510 528-1263